arXiv:1607.06442v1 [cs.DS] 21 Jul 2016

Metric Perturbation Resilience

Konstantin Makarychev
Microsoft Research

Yury Makarychev*
TTIC

Abstract

We study the notion of perturbation resilience introduced by Bilu and Linial (2010) and Awasthi, Blum, and Sheffet (2012). A clustering problem is α-perturbation resilient if the optimal clustering does not change when we perturb all distances by a factor of at most α. We consider a class of clustering problems with center-based objectives, which includes such problems as k-means, k-median, and k-center, and give an exact algorithm for clustering 2-perturbation resilient instances. Our result improves upon the result of Balcan and Liang (2016), who gave an algorithm for clustering $1+\sqrt{2} \approx 2.41$ perturbation resilient instances. Our result is tight in the sense that no polynomial-time algorithm can solve $(2-\varepsilon)$-perturbation resilient instances unless $NP = RP$, as was shown by Balcan, Haghtalab, and White (2016). We show that the algorithm works on instances satisfying a slightly weaker and more natural condition than perturbation resilience, which we call metric perturbation resilience.

*Supported by NSF CAREER award CCF-1150062 and NSF award IIS-1302662.

1 Introduction

In this paper, we present an exact algorithm for solving 2-perturbation resilient instances of clustering problems with natural center-based objectives. The notion of perturbation resilience was proposed by Bilu and Linial [5] and Awasthi, Blum, and Sheffet [1]. Informally, an instance is perturbation resilient if the optimal solution does not change when we perturb the instance. The definition was introduced in the context of beyond-the-worst-case analysis and aims to capture a wide class of real-life instances that are computationally easier than worst-case instances. As several authors argue, in instances arising in practice, the optimal solution is often significantly better than all other solutions and thus does not change if we slightly perturb the instance [5, 4]. It was shown that perturbation resilient instances of such problems as k-center, k-means, k-median, clustering problems with center-based and min-sum objectives, Max Cut, Minimum Multiway Cut, and TSP (with a sufficiently large value of the resilience parameter) can be solved exactly in polynomial-time [1, 2, 3, 4, 5, 6, 7]; while the worst-case instances of these problems are NP-hard.

In a clustering problem, we are given a metric space (X, d) and an integer parameter k; our goal is to partition X into k clusters $C_1, \dots, C_k$ so as to minimize the objective function $\mathcal{H}(C_1, \dots, C_k; d)$ (which depends on the problem at hand). The most well-studied and, perhaps, most interesting clustering objectives are k-means, k-median, and k-center. These objectives are defined as follows. Given a clustering $C_1, \dots, C_k$, the objective is equal to the minimum over all choices of centers $c_1 \in C_1, \dots, c_k \in C_k$ of the following functions:

$$\mathcal{H}_{means}(C_1, \dots, C_k; d) = \sum_{i=1}^{k} \sum_{u \in C_i} d(u, c_i)^2;$$

$$\mathcal{H}_{median}(C_1, \dots, C_k; d) = \sum_{i=1}^{k} \sum_{u \in C_i} d(u, c_i);$$

$$\mathcal{H}_{center}(C_1, \dots, C_k; d) = \max_{i \in \{1,\dots,k\}} \max_{u \in C_i} d(u, c_i).$$

Note that in the optimal solution each cluster C_i consists of those vertices u that are closer to c_i than to other centers c_j; i.e. $(C_1, \dots, C_k)$ is the Voronoi partition of X with centers $c_1, \dots, c_k$. We refer to objectives satisfying this property as *center-based* objectives. We study two closely related classes of center-based objectives – *separable* center-based objectives and *natural* center-based objectives (which we discuss below and formally define in Section 2). We note that k-means, k-median, and k-center are separable and natural center-based objectives. Now we formally define perturbation resilience.

Definition 1.1. *Consider an instance $\mathcal{I} = ((X, d), \mathcal{H}, k)$ of a clustering problem on a metric space (X, d) with objective $\mathcal{H}$. An instance $((X, d'), \mathcal{H}, k)$ is an α-perturbation of $\mathcal{I}$ if*

$$d(u, v) \le d'(u, v) \le \alpha d(u, v) \qquad \text{for every } u, v \in X;$$

here, d' does not have to be a metric. An instance $\mathcal{I}$ is α-perturbation resilient if every α-perturbation of $\mathcal{I}$ has the same optimal clustering as $\mathcal{I}$. We will refer to α as the perturbation resilience parameter.

This definition does not require that the perturbation d' is a metric (d' does not have to satisfy the triangle inequality). It is more natural to consider only *metric* perturbations of $\mathcal{I}$ — those perturbations in which d' is a metric. In this paper, we give the definition of metric perturbation resilience, in which we do require that d' is a metric (see Definition 2.4). (Every α-perturbation resilient instance is also α-*metric* perturbation resilient.)

Known results. Awasthi, Blum, and Sheffet [1] initiated the study of perturbation resilient instances of clustering problems. They offered the definition of a separable center-based objective (s.c.b.o.) and introduced an important center proximity property (see Definition 2.5). They presented an exact algorithm for solving 3-perturbation resilient instances of clustering problems with s.c.b.o.; they also gave an algorithm for $(2+\sqrt{3})$-perturbation resilient instances of clustering problems with s.c.b.o. that have Steiner points. Additionally, they showed that α-perturbation resilient instances of k-median with Steiner points are NP-hard when $\alpha < 3$. Later, Balcan and Liang [3] designed an exact algorithm for $(1+\sqrt{2})$-perturbation resilient instances of problems with s.c.b.o., improving the result of Awasthi, Blum, and Sheffet. Balcan and Liang also studied clustering with the min-sum objective and (α, ε)-perturbation resilience (a weaker notion of perturbation resilience, which we do not discuss in this paper). Recently, Balcan, Haghtalab, and White [2] designed an algorithm for 2-perturbation resilient instances of symmetric and asymmetric k-center and showed that there is no polynomial-time algorithm for $(2-\varepsilon)$ perturbation resilience instances unless $NP = RP$. They also gave an algorithm for 2-perturbation resilient instances of problems with s.c.b.o. satisfying a strong additional condition of cluster verifiability. To summarize, the best known algorithm for arbitrary s.c.b.o. requires that the instance be $1+\sqrt{2} \approx 2.4142$ perturbation resilient [3]; the best known algorithm for k-center requires that the instance be 2-perturbation resilient, and this result cannot be improved [2].

We refer the reader to [6] for an overview of known results for stable instances of combinatorial optimization problems.

Our results. We define a class of clustering problems with *natural center-based objectives*; an objective is a natural center-based objective if it is representable in the following form. For some functions f_c and $g_u(r)$ (f_c is a function of c, $g_u(r)$ is a function of u and r; intuitively, f_{c_i} is the cost of having a center at c_i and $g_u(r)$ is the cost of connecting u to a center at distance r from u), we have

$$\mathcal{H}(C_1, \ldots, C_k; d) = \min_{c_i \in C_i} \sum_{i=1}^{k} \Big(f_{c_i} + \sum_{u \in C_i} g_u(d(u, c_i)) \Big),$$

or

$$\mathcal{H}(C_1, \ldots, C_k; d) = \min_{c_i \in C_i} \max \Big(\max_{i \in \{1,\ldots,k\}} f_{c_i}, \max_{\substack{i \in \{1,\ldots,k\} \\ u \in C_i}} g_u(d(u, c_i)) \Big).$$

This class includes such problems as k-center, k-means, k-median, and metric facility location (for the first three problems: $f_c = 0$ and $g_u(r)$ equals r^2, r, and r respectively).

We present a polynomial-time algorithm for 2-metric perturbation resilient instances of clustering problems with natural center based objectives; thus, we improve the known requirement on the perturbation resilience parameter α from $\alpha \geq 1+\sqrt{2} \approx 2.4142$ to $\alpha \geq 2$ and relax the condition on instances from a stronger α-perturbation resilience condition to a weaker and more natural α-*metric* perturbation resilience condition. Our result is optimal, since even $(2-\varepsilon)$-perturbation resilient instances of k-center cannot be solved in polynomial time unless $NP = RP$ [2]. In particular, our result improves the requirement for k-median and k-means from $\alpha \geq 1+\sqrt{2}$ to $\alpha \geq 2$.

Theorem 1.2. *There exists a polynomial-time algorithm that given any 2-metric perturbation resilient instance $((X, d), \mathcal{H}, k)$ of a clustering problem with natural center-based objective, returns the (exact) optimal clustering of X.*

Our algorithm is quite simple. It first runs the single-linkage algorithm to construct the minimum spanning tree on points of X and then partitions the minimum spanning tree into k clusters

using dynamic programming. We note that Awasthi, Blum, and Sheffet [1] also used the single-linkage algorithm together with dynamic programming to cluster 3-perturbation resilient instances. However, their approach is substantially different from ours: They first find a hierarchical clustering of X using the single-linkage algorithm and then pick k optimal clusters from this hierarchical clustering. This approach fails for α-perturbation resilient instances with $\alpha < 3$ (see [1]). That is why, we do not use the single-linkage *hierarchical* clustering in our algorithm, and, instead, partition the minimum spanning tree.

We note that the definitions of *separable* and *natural* center-based objectives are different. However, in Section 5 we consider a slightly strengthened definition of s.c.b.o. and show that every s.c.b.o., under this new definition, is also a natural center-based objective; thus, our result applies to it. We are not aware of any non-pathological objective that satisfies the definition of s.c.b.o. but is not a natural center-based objective.

Finally, we consider clustering with s.c.b.o. and show that the optimal solution for every α-metric perturbation resilient instance satisfies the α-center proximity property; previously, that was only known for α-perturbation resilient instances [1]. Our result implies that the algorithms by Balcan and Liang [3] and Balcan, Haghtalab, and White [2] for clustering with s.c.b.o. and k-center, respectively, apply not only to α-perturbation resilient but also α-metric perturbation resilient instances.

Overview. In Section 2, we formally define key notions used in this paper. Then, in Section 3, we prove that the optimal solution for every α-metric perturbation resilient instance satisfies the α-center proximity property. We use this result later in the analysis of our algorithm; also, as we noted above, it is of independent interest and implies that previously known algorithms from [2, 3] work under the metric permutation resilience assumption. In Section 4, we present our algorithm for solving 2-perturbation resilient instances of problems with natural center-based objectives. The algorithm consists of two steps. In the first step, we find a minimum spanning tree in the metric space (X, d). In the second step, we partition the tree into k subtrees using dynamic programming and then output the corresponding clustering of X. Finally, in Section 5, we prove that if a s.c.b.o. satisfies some additional properties, then it is a natural center-based objective.

2 Preliminaries

In this section, we formally define key notions used in this paper: a clustering problem, α-metric perturbation resilience, separable center-based and natural center-based objectives.

Definition 2.1. *An instance of a clustering problem is the tuple $((X, d), \mathcal{H}, k)$ of a metric space (X, d), objective function $\mathcal{H}$, and integer number $k > 1$. The objective $\mathcal{H}$ is a function that given a partition of X into k sets $C_1, \ldots, C_k$ and a metric d on X returns a nonnegative real number, which we call the cost of the partition.*

Given an instance of a clustering problem $((X, d), \mathcal{H}, k)$, our goal is to partition X into disjoint sets $C_1, \ldots, C_k$, so as to minimize $\mathcal{H}(C_1, \ldots, C_k; d)$. Now we define center-based and separable center-based objectives. We note that the definitions of Awasthi et al. [1] make several implicit assumptions that we make explicit here.

Definition 2.2. *We say that $\mathcal{H}$ is a center-based objective function if the following properties hold.*

1. *Given a subset $S \subset X$ and distance d_S on S, we can find the optimal center $c \in S$ for S, or, if there is more than one choice of an optimal center, a set of optimal centers* $\mathrm{center}(S, d_S) \subset S$. *(In the former case,* $\mathrm{center}(S, d_S) = \{c\}$*).*

2. *For every metric d on X, there exists an optimal clustering $C_1, \dots, C_k$ of X (i.e., a clustering that minimizes $\mathcal{H}(C_1, \dots, C_k; d)$) such that for every two distinct clusters C_i, C_j, every point p in C_i and centers $c_i \in \mathrm{center}(C_i, d|_{C_i})$ and $c_j \in \mathrm{center}(C_j, d|_{C_j})$ we have $d(p, c_i) \le d(p, c_j)$. Furthermore, if the optimal clustering $C_1, \dots, C_k$ is unique, then $d(p, c_i) < d(p, c_j)$.*

The objective is separable if, additionally, we can define individual cluster scores so that the following holds.

- *The cost of the clustering is either the sum (for separable sum-objectives) or maximum (for separable max-objectives) of the cluster scores.*
- *The score $H(C, d|_C)$ of each cluster C depends only on C and $d|_C$, and can be computed in polynomial time.*

In this paper, we consider a slightly narrower class of *natural* center-based objectives (which we described in the introduction). The class contains most important center-based objectives: k-center, k-means, and k-median, as well as the facility location objective. We are not aware of any reasonable center-based objective that is not a natural center-based objective. Now, we formally define natural center-based objectives.

Definition 2.3. *We say that $\mathcal{H}$ is a natural center-based objective function for a ground set X, if there exist functions $f : X \to \mathbb{R}$ and $g : X \times \mathbb{R} \to \mathbb{R}$ such that*

$$\mathcal{H}(C_1, \dots, C_k; d) = \min_{c_i \in C_i} \sum_{i=1}^{k} \Big(f_{c_i} + \sum_{u \in C_i} g_u(d(u, c_i)) \Big), \tag{1}$$

or

$$\mathcal{H}(C_1, \dots, C_k; d) = \min_{c_i \in C_i} \max \Big(\max_{i \in \{1,\dots,k\}} f_{c_i}, \max_{\substack{i \in \{1,\dots,k\} \\ u \in C_i}} g_u(d(u, c_i)) \Big). \tag{2}$$

We require that the functions f and g be computable in polynomial time, and that g_u be non-decreasing for every $u \in X$. We call points c_i the centers of the clustering.

We could have also defined an objective with an ℓ_p-aggregate function $\mathcal{H} = \big(\sum_{i=1}^{k} \big(f_{c_i}^p + \sum_{u \in C_i} g_u(d(u, c_i))^p \big)^{1/p}$, but it is equivalent to (1) with $f'(c) = f^p(c)$ and $g'(u, r) = g^p(u, r)$.

Now, we formally define a metric perturbation and metric perturbation resilience. Since we do not require that the objective $\mathcal{H}$ is homogeneous as a function of the metric d, we introduce two perturbation resilience parameters α_1 and α_2 in the definition, which specify by how much the distances can be contracted and expanded, respectively, in the perturbed instances.

Definition 2.4. *Consider a metric space (X, d). We say that a metric d' is an (α_1, α_2)-metric perturbation of (X, d) if $\alpha_1^{-1} d(u, v) \le d'(u, v) \le \alpha_2 d(u, v)$ for every $u, v \in X$. An instance $((X, d), \mathcal{H}, k)$ is (α_1, α_2)-metric perturbation resilient if for every (α_1, α_2)-metric perturbation d' of d, the unique optimal clustering for $((X, d'), \mathcal{H}, k)$ is the same as for $((X, d'), \mathcal{H}, k)$.*

Note that the centers of clusters in the optimal solutions for $((X, d), \mathcal{H}, k)$ and $((X, d'), \mathcal{H}, k)$ may differ. We say that an instance $(\mathcal{H}, (X, d), k)$ is α-metric perturbation resilient if it is $(\alpha, 1)$-metric perturbation resilient.

Observe that if instance $((X, d), \mathcal{H}, k)$ is (α_1, α_2)-metric perturbation resilient, then $((X, \lambda d), \mathcal{H}, k)$ is $(\lambda\alpha_1, \lambda\alpha_2)$-metric perturbation resilient for $\lambda \in [\alpha_1^{-1}, \alpha_2]$. Particularly, if $((X, d), \mathcal{H}, k)$ is (α_1, α_2)-metric perturbation resilient, then $((X, \alpha_2 d), \mathcal{H}, k)$ is $(\alpha, 1)$-metric perturbation resilient and the

optimal solution for $((X, \alpha_2 d), \mathcal{H}, k)$ is the same as for $((X, d), \mathcal{H}, k)$. Thus, to solve an (α_1, α_2)-metric perturbation resilient instance $((X, \lambda d), \mathcal{H}, k)$, it suffices to solve $\alpha = (\alpha_1 \alpha_2)$ metric perturbation resilient instance $((X, \alpha_2 d), \mathcal{H}, k)$. Consequently, we will only consider α-metric perturbation resilient instances in this paper.

Finally, we recall the definition of the α-center proximity property, introduced in [1].

Definition 2.5. *We say that a clustering $C_1, \ldots, C_k$ of X with centers $c_1, \ldots, c_k$ satisfies the α-center proximity property if for all $i \neq j$ and $p \in C_i$, we have $d(p, c_j) > \alpha d(p, c_i)$.*

3 Center Proximity for Metric Perturbation Resilience

In this section, we prove that the (unique) optimal solution to an *α-metric* perturbation resilient clustering problem satisfies the α-center proximity property. Our proof is similar to the proof of Awasthi, Blum, and Sheffet, who showed that the optimal solution to a (non-metric) α-perturbation resilient clustering problem satisfies the α-center proximity property.

Theorem 3.1. *Consider an α-metric perturbation resilient clustering problem $((X, d), \mathcal{H}, k)$ with a center-based objective. Let $C_1, \ldots, C_k$ be the unique optimal solution; and let $c_1, \ldots, c_k$ be a set of centers of $C_1, \ldots, C_k$ (that is, each $c_i \in \text{center}(C_i, d|_{C_i})$). Then, the following α-center proximity property holds: for all $i \neq j$ and $p \in C_i$, we have $d(p, c_j) > \alpha d(p, c_i)$.*

Proof. Suppose that $d(p, c_j) \leq \alpha d(p, c_i)$. Let $r^* = d(p, c_i)$. Define a new metric d' as follows. Consider the complete graph on X. Assign length $len(u, v) = d(u, v)$ to each edge (u, v) other than (p, c_j). Assign length $len(p, c_j) = r^*$ to the edge (p, c_j). Let metric $d'(u, v)$ be the shortest path metric on the complete graph on X with edge lengths $len(u, v)$. Note that $d(p, c_j) \geq d(p, c_i) = r^*$ since $p \in C_i$ and $C_1, \ldots, C_k$ is an optimal clustering. Hence, for every (u, v): $len(u, v) \leq d(u, v)$ and $d'(u, v) \leq d(u, v)$. It is easy to see that

$$d'(u, v) = \min(d(u, v), d(u, p) + r^* + d(c_j, v), d(v, p) + r^* + d(c_j, u)).$$

Observe that since the ratio $d(u, v)/len(u, v)$ is at most $d(p, c_j)/r^* \leq \alpha$ for all edges (u, v), we have $d(u, v)/d'(u, v) \leq \alpha$ for all u and v. Hence, $d(u, v) \leq \alpha d'(u, v) \leq \alpha d(u, v)$, and consequently, d' is an $(\alpha, 1)$-metric perturbation of d.

We now show that d' is equal to d within the cluster C_i and within the cluster C_j.

Lemma 3.2. *For all $u, v \in C_i$, we have $d(u, v) = d'(u, v)$, and for all $u, v \in C_j$, we have $d(u, v) = d'(u, v)$.*

Proof. I. Consider two points u, v in C_i. We need to show that $d(u, v) = d'(u, v)$. It suffices to prove that

$$d(u, v) \leq \min(d(u, p) + r^* + d(c_j, v), d(v, p) + r^* + d(c_j, u)).$$

Assume without loss of generality that $d(u, p) + r^* + d(c_j, v) \leq d(v, p) + r^* + d(c_j, u)$. We have

$$d(u, p) + r^* + d(c_j, v) = d(u, p) + d(p, c_i) + d(c_j, v) \geq d(u, c_i) + d(c_j, v).$$

Since $v \in C_i$, we have $d(v, c_i) \leq d(v, c_j)$, and thus

$$d(u, p) + r^* + d(c_j, v) \geq d(u, c_i) + d(c_i, v) \geq d(u, v).$$

II. Consider two points u, v in C_j. Similarly to the previous case, we need to show that $d(u,v) \le d(u,p) + r^* + d(c_j,v)$. Since now $u \in C_j$, we have $d(u,c_j) \le d(u,c_i)$. Thus,

$$d(u,p)+r^*+d(c_j,v) = \big(d(u,p)+d(p,c_i)\big)+d(c_j,v) \ge d(u,c_i)+d(c_j,v) \ge d(u,c_j)+d(c_j,v) \ge d(u,v).$$

□

By the definition of α-metric perturbation stability, the optimal clusterings for metrics d and d' are the same. By Lemma 3.2, the distance functions d and d' are equal within the clusters C_i and C_j. Hence, the centers of C_i and C_j w.r.t. metric d' are also points c_i and c_j, respectively (see Definition 2.2, item 1). Thus, by the definition of center-based objective, $d'(c_i,p) < d'(c_j,p)$, and, consequently,

$$d(c_i,p) = d'(c_i,p) < d'(c_j,p) = r^* = d(c_i,p).$$

We get a contradiction, which finishes the proof. □

Corollary 3.3. *Consider a 2-metric perturbation resilient instance. Let $C_1,\dots,C_k$ be an optimal clustering with centers $c_1,\dots,c_k$. Then each point u in C_i is closer to c_i than to any point v outside of C_i.*

Proof. Suppose that $v \in C_j$ for some $j \ne i$. By the triangle inequality, $d(u,v) \ge d(u,c_j) - d(c_j,v)$. By Lemma 3.1, $d(u,c_j) > 2d(u,c_i)$ and $d(c_j,v) < d(c_i,v)/2$. Thus,

$$d(u,v) \ge d(u,c_j) - d(c_j,v) > 2d(u,c_i) - \frac{d(c_i,v)}{2} \ge 2d(u,c_i) - \frac{d(c_i,u)+d(u,v)}{2}.$$

In the last inequality, we used the triangle inequality $d(c_i,v) \le d(c_i,u) + d(u,v)$. Rearranging the terms, we get $d(u,v) > d(c_i,u)$. □

4 Clustering Algorithm

In this section, we present our algorithm for solving 2-metric perturbation resilient instances of clustering problems with natural center-based objectives. Our algorithm is based on single–linkage clustering: first, we find a minimum spanning tree on the metric space X (e.g., using Kruskal's algorithm) and then run a dynamic programming algorithm on the spanning tree to find the clusters. We describe the two steps of the algorithm in Sections 4.1 and 4.2.

4.1 Minimum Spanning Tree

At the first phase of the algorithm, we construct a minimum spanning tree on the points of the metric space using Kruskal's algorithm. Kruskal's algorithm maintains a collection of trees. Initially, each tree is a singleton point. At every step, the algorithm finds two points closest to each other that belong to different trees and adds an edges between them. The algorithm terminates when all points belong to the same tree. We denote the obtained spanning tree by T. The key observation is that each cluster C_i forms a subtree of the spanning tree T.

Lemma 4.1. *Each cluster C_i forms a subtree of the spanning tree T. In other words, the unique path between every two vertices $u, v \in C_i$ does not leave the cluster C_i.*

Proof. Let c_i be the center of C_i. We show that the (unique) path p from u to c_i lies in C_i, and, therefore, the statement of Lemma 4.1 holds. Let u' be the next vertex after u on the path p. Consider the step at which Kruskal's algorithm added the edge (u, u'). At that step, u and c_i were in distinct connected components (as p is the only path connecting u and c_i). Thus, $d(u, u') \leq d(u, c_i)$ as otherwise the algorithm would have added the edge (u, c_i) instead of (u, u'). By Corollary 3.3, the inequality $d(u, u') \leq d(u, c_i)$ implies that u' belongs to C_i. Proceeding by induction we conclude that all vertices on the path p belong to C_i. □

4.2 Dynamic Programming Algorithm

At the second phase, we use dynamic programming to compute the optimal clustering. We root the tree T at an arbitrary vertex. We denote its root by *root* and the subtree rooted at u by T_u. We first assume that the tree is binary. Later, we explain how to transform any tree into a binary tree by adding dummy vertices.

The algorithm partitions the tree into (non-empty) subtrees $P_1, \ldots, P_k$ and assigns a center $c_i \in P_i$ to all vertices in the subtree P_i so as to minimize the objective:

$$\sum_{i=1}^{k} f_{c_i} + \sum_{i=1}^{k} \sum_{u \in P_i} g_u(d(u, c_i)). \tag{3}$$

Lemma 4.1 implies that the optimal partitioning of X is the solution to this problem.

Let $cost_u(k', c)$ be the minimum cost (3) of partitioning the subtree T_u into k' subtrees $P_1, \ldots, P_k$ and choosing k centers $c_1, \ldots, c_k$ so that the following conditions hold:

1. $u \in P_1$ and $c_1 = c$ (we denote the tree that contains u by P_1 and require that its center be c).
2. if $c_1 \in T_u$, then $c_1 \in P_1$ (if the center c_1 of P_1 lies in T_u, then it must be in P_1);
3. $c_i \in P_i$ for $i > 1$ (the center c_i of every other tree P_i lies in P_i).

That is, we assume that u belongs to the first subtree P_1 and that c is the center for P_1. Every center c_i must belong to the corresponding set P_i except for c_1. However, if $c_1 \in T_u$, then $c_1 \in P_1$.

Denote the children of vertex u by l_u and r_u (recall that we assume that the tree is binary). The cost $\text{cost}_u(k', c)$ is computed using the following recursive formulas: if $c \notin T_{l_u} \cup T_{r_u}$, then

$$\begin{aligned}
\text{cost}_u(j, c) = & f_c + g_u(d(c, u)) + \min\big(&& (4)\\
& \min\{\text{cost}_{l_u}(j', c') + \text{cost}_{r_u}(j'', c'') : j' + j'' = j - 1,\ c' \in T_{l_u}, c'' \in T_{r_u}\}, && (5)\\
& \min\{\text{cost}_{l_u}(j', c') + \text{cost}_{r_u}(j'', c) - f_c : j' + j'' = j,\ c' \in T_{l_u}\}, && (6)\\
& \min\{\text{cost}_{l_u}(j', c) + \text{cost}_{r_u}(j'', c'') - f_c : j' + j'' = j,\ c'' \in T_{r_u}\}, && (7)\\
& \min\{\text{cost}_{l_u}(j', c) + \text{cost}_{r_u}(j'', c) - 2f_c : j' + j'' = j + 1\}\big). && (8)
\end{aligned}$$

If $c \in T_{l_u}$, then we remove lines (5) and (6) from the formula. If $c \in T_{r_u}$, then we remove lines (5) and (7) from the formula.

The first term $f_c + g_u(d(c, u))$ is the cost of opening a center in c and assigning u to c. The lines (5–8) correspond to the following cases:

(5) neither l_u nor r_u is in P_1; they are assigned to (trees P_i and P_j with centers) c' and c'',

(6) l_u is not in P_1, but r_u is in P_1; they are assigned to c' and c,

(7) l_u is in P_1, but r_u is not in P_1; they are assigned to c and c'',

(8) both l_u and r_u are in P_1; they are assigned to c.

For leaves, we set $\text{cost}_u(1, c) = f_c + g_u(d(u, c))$ and $\text{cost}_u(j, c) = \infty$ for $j > 1$. Note that if we want to find a partitioning of T into at most k subtrees, we can use a slightly simpler dynamic program. It is easy to verify that the formulas above hold. The cost of the optimal partitioning of T into k subtrees equals $\min_{c \in X} \text{cost}_{root}(k, c)$.

We now explain how to transform the tree T into a binary tree. If a vertex u has more than two children, we add new dummy vertices between u and its children by repeating the following procedure: take a vertex u having more than two children; pick any two children of u: u_1 and u_2; create a new child v of u and rehang subtrees T_{u_1} and T_{u_2} to the vertex v. We forbid opening centers in dummy vertices v by setting the opening cost f_v to be infinity. We set the assignment costs g_v to be 0. Note that Lemma 4.1 still holds for the new tree if we place every dummy vertex in the same part P_i as its parent.

5 Universal Separable Center-Based Clustering Objectives

In this section, we show that every separable center-based objective, satisfying some additional properties, is a natural center-based objective. To this end, we define a *universal* center-based objective and show that every universal center-based objective is a natural center-based objective.

Loosely speaking, a universal center-based objective is a center-based objective that satisfies two properties, which we will discuss now:

- An arbitrary center-based objective is defined on a specific set of points X and can be used to compute the cost of clustering only of the set X (given a metric d on X). **In contrast, a *universal objective* can be used to compute the cost of clustering of *any* ground set X.**

- Recall that in every optimal clustering with a center-based objective each point $u \in X$ is closer to the center of its own cluster than to the center of any other cluster. If a partition is not optimal, some points might be closer to the centers of clusters that do not contain them than to the centers of their own clusters. Then, we can move such points to other clusters so as to minimize their distance to the cluster centers. In fact, one of the two steps of Lloyd's algorithm does exactly this; hence, we call such a transformation *a Lloyd's improvement*. We slightly strengthen the definition of a center-based objective by requiring that **a *universal objective* not increase when we make a Lloyd's improvement of *any* – not necessarily optimal – clustering**.

Now we give a few auxiliary definitions and then formally define universal center-based objectives. Since data sets used in applications are usually labeled, we will consider "labeled metric spaces". We will assume that the cost of clustering of X may depend on the distances between the points in X and point labels (but not on the identities of points).

Definition 5.1 (Labeled Metric Space)**.** *A metric space labeled with a set of labels L is a pair $((X, d), l)$, where (X, d) is a metric space, and $l : X \to L$ is a function that assigns a label to each point in X.*

Definition 5.2 (Isomorphic Labeled Metric Spaces)**.** *We say that two metric spaces $((X', d'), l')$ and $((X'', d''), l'')$ labeled with the same set L are isomorphic if there exists an isometry $\varphi : X' \to X''$ (i.e., φ is a bijection that preserves distances: $d'(u, v) = d''(\varphi(u), \varphi(v))$ for all $u, v \in X'$) that preserves labels; i.e. $l'(u) = l''(\varphi(u))$ for all $u \in X'$.*

We note that in the definition above the set L may be infinite. We denote the restriction of $((X,d),l)$ to a non-empty subset $C \subset X$ by $((X,d),l)|_C$: $((X,d),l)|_C = ((C,d|_C),l|_C)$. Note that the restriction of a metric set labeled with L to a cluster C is also a metric set labeled with L.

Definition 5.3. *Consider a clustering problem $((X,d),\mathcal{H},k)$ with a center-based objective. We say that a clustering $C'_1,\dots,C'_k$ is a Lloyd's improvement of a clustering $C_1,\dots,C_k$ if there exists a set of centers $c_1,\dots,c_k$ of $C_1,\dots,C_k$ (i.e., each $c_i \in \mathrm{center}(C_i,d)$) such that*

- *$c_i \in C'_i$ (a Lloyd's improvement does not move the centers to other clusters)*
- *for every $x \in X$: if $x \in C_i$ and $x \in C'_j$, then $d(x,c_j) \le d(x,c_i)$ (a Lloyd's improvement may move point x from C_i to C'_j only if $d(c,c_j) \le d(x,c_i)$).*

Definition 5.4 (Universal Objective)**.** *We say that $\mathcal{H}$ is a universal center-based clustering objective for a label set L, if for every metric space $((X,d),l)$ labeled with L the problem $((X,d),\mathcal{H}_l,k)$ is a clustering problem with a separable center-based objective and the following two conditions hold.*

1. *Cluster scores H_l are universal ("can be used on any metric space"): Given any finite metric space $((C,d),l)$ labeled with L, the function $H_l(C,d)$ returns a real number – the cost of C; and $H_{l'}(C',d') = H_{l''}(C'',d'')$ for any two isomorphic labeled metric spaces $((C',d'),l')$ and $((C'',d''),l'')$.*
2. *If $C'_1,\dots,C'_k$ is a Lloyd's improvement of $C_1,\dots,C_k$, then $\mathcal{H}(C'_1,\dots,C'_k;d) \le \mathcal{H}(C_1,\dots,C_k;d)$.*

Note that every natural center-based objective is a universal objective. The label set is the set of pairs (f,g), where $f \in \mathbb{R}$ is a real number; g is a nondecreasing function from $\mathbb{R}^{\ge 0}$ to $\mathbb{R}$. Every point $x \in X$ is assigned the label $l(x) = (f_x, h_x)$. The score of a cluster C equals

$$H_l(C,d) = \min_{c\in C}\Big(f_c + \sum_{x\in C} g_x(d(x,c))\Big); \tag{9}$$

$$H_l(C,d) = \min_{c\in C}\Big(\max\big(f_c, \max_{x\in C} g_x(d(x,c))\big)\Big). \tag{10}$$

It is easy to see that Lloyd's improvements may only decrease the cost of a clustering, since the functions g_x are non-decreasing. We now show that every clustering problem with universal separable center-based objectives is a problem with natural center-based objectives.

Theorem 5.5. *I. Let $((X,d),\mathcal{H},k)$ be a clustering problem with a universal center-based separable sum-objective. Then, the scoring function H_l can be represented as (9) for some nondecreasing functions f and g such that the minimum is attained when c is a center of C; and thus $\mathcal{H}(C_1,\dots,C_k;d) = \sum_{i=1}^k H_l(C_i;d|_{C_i})$ is a natural center-based objective.*

II. Let $((X,d),\mathcal{H},k)$ be a clustering problem with a universal center-based separable max-objective. If the cost of any singleton cluster $\{x\}$ equals 0, then the scoring function H_l can be represented as (10) for some nondecreasing functions f and g such that the minimum is attained when c is a center of C; ; and thus $\mathcal{H}(C_1,\dots,C_k;d) = \max_{i\in\{1,\dots,k\}} H_l(C_i;d|_{C_i})$ is a natural center-based objective.

Proof. I. For every label $a \in L$ define two special labeled metric spaces M_a and $M_{a,a,r}$. The metric space M_a is the metric space on a single point labeled with a; the metric space $M_{a,a,r}$ is the metric space on two points at distance r that are both labeled with a. Consider a point $x \in X$ labeled with a. Let $f_x = H(M_a)$ and $g_x(r) = H(M_{a,a,r}) - f_x$. We claim that the cost of any cluster C equals

$$\min_{c\in C}\big(f_c + \sum_{x\in C\setminus\{c\}} g_x(d(x,c))\big),$$

and the minimum is attained when c is a center of C. Consequently, equation (9) holds for functions $\tilde{f}_c = f_c - g_c(0)$ and g_x. To prove the claim, we construct a new metric space C'_p for every $p \in C$. The metric space C'_p consists of all points of C and new points x', which are "duplicates" of points $x \in C \setminus \{p\}$. We do not change distances between points in C and labels on the points from C. We copy the label of each x to x'; i.e., we let $l(x') = l(x)$. We place x' at the same location as the point p. That is, we let $d(x', y) = d(p, y)$ for all $x, y \in C$. Consider two clusterings of C'_p. The first clustering partitions C'_p into the cluster C and singleton clusters $\{x'\}$ (there is a singleton cluster for each $x \in C \setminus \{p\}$). The cost of each singleton cluster $\{x'\}$ equals f_x, since $\{x'\}$ is isomorphic to $M_{l(x)}$. Hence, the cost of the first clustering equals

$$H_l(C, d|_C) + \sum_{x \in C \setminus \{p\}} f_x. \tag{11}$$

The second clustering partitions C' into pairs of points $\{x, x'\}$ for each $x \in C \setminus \{p\}$ and the singleton cluster $\{p\}$. The cost of each cluster $\{x, x'\}$ equals

$$H(M_{a,a,d(x,x')}) = g_x(d(x, x')) + f_x = g_x(d(x, p)) + f_x,$$

since the cluster $\{x, x'\}$ is isomorphic to $M_{a,a,d(x,x')}$. The cost of $\{p\}$ equals f_p. Hence, the cost of the second clustering equals

$$f_p + \sum_{x \in C \setminus \{p\}} \big(f_x + g_x(d(x, c))\big). \tag{12}$$

Note that the points x and x' in each cluster $\{x, x'\}$ are symmetric, so we may assume that x' is the center of the cluster $\{x, x'\}$. Observe that the first clustering is a Lloyd's improvement of the second clustering: To get the first clustering from the second clustering we move x from each cluster $\{x, x'\}$ to the singleton cluster $\{p\}$. We are allowed to do that since $d(x, p) = d(x, x')$. Hence, the cost of the second clustering is upper bounded by the cost of the first clustering. We have from (11) and (12),

$$H_l(C, d|_C) \le f_p + \big(\sum_{x \in C \setminus \{p\}} f_x + g_x(d(x, c))\big) - \sum_{x \in C \setminus \{p\}} f_x = f_p + \sum_{x \in C \setminus \{p\}} g_x(d(x, c)). \tag{13}$$

On the other hand, if p is the center of C, then the second clustering is a Lloyd's improvement of the first one. Hence, (13) is an equality when p is the center of C.

It remains to show that functions g_x are non-decreasing. We need to prove that for every $r_1 \ge r_2 \ge 0$, $g_x(r_1) \ge g_x(r_2)$. We create a metric space on three points x', x'', x''' each having the same label as x. Let $d(x', x'') = r_1$, $d(x'', x''') = r_2$, $d(x', x''') = r_1 + r_2$. The cost of clustering $\{x', x''\}, \{x'''\}$ equals $2f_x + g_x(r_1)$; the cost of clustering $\{x'\}, \{x'', x'''\}$ equals $2f_x + g_x(r_2)$. Since $r_1 \ge r_2$, the second clustering is a Lloyd's improvement of the first one. Hence, $2f_x + g_x(r_2) \le 2f_x + g_x(r_1)$ and $g_x(r_2) \le g_x(r_1)$. This completes the proof of part I. The proof of part II is similar. □

References

[1] Pranjal Awasthi, Avrim Blum, and Or Sheffet. Center-based clustering under perturbation stability. *Information Processing Letters*, 112(1):49–54, 2012.

[2] Maria-Florina Balcan, Nika Haghtalab, and Colin White. Symmetric and asymmetric k-center clustering under stability. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*, 2016.

[3] Maria-Florina Balcan and Yingyu Liang. Clustering under perturbation resilience. *SIAM Journal on Computing*, 2016.

[4] Yonatan Bilu, Amit Daniely, Nati Linial, and Michael Saks. On the practically interesting instances of maxcut. In *Proceedings of the Symposium on Theoretical Aspects of Computer Science*, pages 526–537, 2013.

[5] Yonatan Bilu and Nathan Linial. Are stable instances easy? In *Innovations in Computer Science*, pages 332–341, 2010.

[6] Konstantin Makarychev, Yury Makarychev, and Aravindan Vijayaraghavan. Bilu—Linial stable instances of Max Cut and Minimum Multiway Cut. In *Proceedings of the Symposium on Discrete Algorithms*, pages 890–906, 2014.

[7] Matúš Mihalák, Marcel Schöngens, Rastislav Šrámek, and Peter Widmayer. On the complexity of the metric tsp under stability considerations. In *SOFSEM 2011: Theory and Practice of Computer Science*, pages 382–393. Springer, 2011.